

SHARED ELECTRONIC INK ANNOTATION METHOD AND SYSTEM

CROSS-REFERENCE TO RELATED APPLICATION

The present invention claims priority to United States
5 provisional patent application Serial No. 60/471,844, filed
May 19, 2003 and herein incorporated by reference.

FIELD OF THE INVENTION

The invention relates generally to computer systems and
10 networks, and more particularly to computer systems employing
electronic inking capabilities.

BACKGROUND OF THE INVENTION

Contemporary technology such as a tablet PC allows
15 individuals to write notes and other information such as
sketches onto their computer screens. Handwriting may be
recognized as text, and documents and the like may have
annotations added thereto by annotation mechanisms.
Annotations may even be searched as if they were text, as
20 generally described United States Patent application Serial
No. 10/367,198, assigned to the assignee of the present
invention. Such electronic searching capabilities provide
benefits that are superior to paper notes.

However, while formal electronic documents may be
25 annotated, in other scenarios existing annotation mechanisms

still fall short of what users can do with paper. For example, such mechanisms are not able to reproduce the experience of several individuals seated around a table writing directly on a shared paper document. Although a
5 computer user could email an annotated electronic document to other users, the step of emailing removes the immediacy that paper users would experience. It would be highly burdensome and interruptive for a user to have to perform such a file transfer every time the user wrote something down.

10 What is needed is a way for computer users to exchange electronic ink annotations with other users, in a manner that is essentially immediate and automatic, and handles one-to-many or many-to-many real time situations, such as note taking during a conference. The annotations should be useable with
15 formal documents, but also should be able to be used with any publication, such as a slide of a presentation that appears on a user's computer screen.

SUMMARY OF THE INVENTION

20 Briefly, the present invention provides a system and method in a computing environment such as a computer network in which computer users are able to make (at least some of) their handwritten annotations public to other computer users, as well as to view notes from other computer users who have

chosen to make their notes public. Users can selectively publish authored notes, and subscribe to other user's published notes. A computer user that has made his or her notes public is also able to determine who is subscribing to
5 those notes.

To provide the system and method, a computer user runs a program such as on a tablet-based personal computer, and enters a publication mode, presented to the user as a public data entry area. In this mode, ink strokes are sent to a
10 server, which records the ink strokes in association with the user's identity, and with the current context of those notes, such as a page (e.g., slide) of a publication (e.g., a PowerPoint presentation). The server also tracks what other users have subscribed to public notes corresponding to that
15 identity.

In one implementation, while a user is in the publishing mode, a background thread executing on the user's computer operates to send the stroke data to the server, essentially as it is being entered. A separate background thread
20 communicates with the server to receive and provide the public stroke data of other users to which the computer user has subscribed.

Various color schemes may be employed to help a user differentiate among the user's own notes and notes from

others. A user may also filter the view of the shared annotations, so as to unclutter the page and/or narrow down authorship of the notes they are viewing.

The present invention also facilitates other useful applications, including a mode that provides a wholly public persistent space which any computer user can annotate, referred to as a shared canvas. For example, such a space is highly useful in an environment having audience members and speakers, such as to make background comments, project this "canvas" for those present to view a spontaneous sketch or diagram, collect points of view for an informal poll, and so forth. The space is tied to a publication page being presented, such as a slide of a slide show presentation, and typically controlled by one individual, generally the one making the presentation. Another mode provides a graffiti-type wall in which users can add any annotations, independent of any page.

Other advantages will become apparent from the following detailed description when taken in conjunction with the drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram generally representing a computer system into which the present invention may be incorporated;

5 FIG. 2 is a block diagram generally representing a network in which shared annotations have been implemented in accordance with an aspect of the present invention;

FIG. 3 is a block diagram generally representing data structures used by the server to track publishers (annotation
10 authors), their public annotations, and subscribers, in accordance with an aspect of the present invention;

FIG. 4 is a representation of a program display showing an example interface for receiving and publishing presentation-related public annotations, in accordance with an
15 aspect of the present invention;

FIG. 5 is a representation of a program display showing an example of how a user can publish notes to a shared canvas, in accordance with an aspect of the present invention; and

FIG. 6 is a representation of a program display showing
20 an example of how a user can publish notes to any other user, including non-subscribers, in accordance with an aspect of the present invention.

DETAILED DESCRIPTION

EXEMPLARY OPERATING ENVIRONMENT

FIGURE 1 illustrates an example of a suitable computing system environment 100 on which the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, tablet devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

The invention may be described in the general context of computer-executable instructions, such as program modules,

being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, and so forth that perform particular tasks or implement particular abstract data types. The invention may
5 also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory
10 storage devices.

With reference to FIG. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of the computer 110 may include, but are not limited to, a
15 processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using
20 any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards

Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

The computer 110 typically includes a variety of computer-readable media. Computer-readable media can be any available media that can be accessed by the computer 110 and includes both volatile and nonvolatile media, and removable and non-removable media. By way of example, and not limitation, computer-readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer 110. Communication media typically embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data

signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a
5 wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above should also be included within the scope of computer-readable media.

The system memory 130 includes computer storage media in
10 the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically
15 stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way of example, and not limitation, FIG. 1 illustrates operating system 134, application programs 135, other program modules
20 136 and program data 137.

The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile

magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

The drives and their associated computer storage media, discussed above and illustrated in FIG. 1, provide storage of computer-readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146 and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and

program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers herein to illustrate that, at a minimum, they are different copies. A user may enter commands and
5 information into the computer 110 through input devices such as a tablet (electronic digitizer) 164, a microphone 163, a keyboard 162 and pointing device 161, commonly referred to as mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish,
10 scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A
15 monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. The monitor 191 may also be integrated with a touch-screen panel 193 or the like that can input digitized input such as handwriting into the computer system 110 via an
20 interface, such as a touch-screen interface 192. Note that the monitor and/or touch screen panel can be physically coupled to a housing in which the computing device 110 is incorporated, such as in a tablet-type personal computer, wherein the touch screen panel 193 essentially serves as the

tablet 164. In addition, computers such as the computing device 110 may also include other peripheral output devices such as speakers 195 and printer 196, which may be connected through an output peripheral interface 194 or the like.

5 The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes
10 many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks.
15 Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170, and particularly may include one configured for
20 wireless networking. When used in a WAN networking environment, the computer 110 may also include a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user

input interface 160 or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not
5 limitation, FIG. 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

10

SHARED ANNOTATIONS

The present invention, in part, directed towards a method and system that use electronic ink, such as formatted in the form of stroke data, to provide annotations to other computer
15 users. While the examples herein generally refer to handwriting capable devices in the form of portable computer systems such as tablet personal computers, it should be understood that virtually any device capable of transmitting and receiving data can participate in the shared annotation
20 schemes described herein. For example, a user could type notes and send text rather than stroke data to other users, and bitmaps or other image data may also be sent, such as to users that do not have computers capable of interpreting stroke data. A translation mechanism or the like may be used

to convert strokes to recognized text and/or bitmaps for such users. Note however that tablet PCs are particularly suited to the present invention, as they provide the advantages of mobility and the benefits of handwritten entry, including more
5 natural input, more expressive comments, and the absence of the sounds of typing. Thus, in a described implementation, ink is sent as stroke vectors to be recreated on individual clients, rather than as bitmap images, generally to reduce network traffic, but in principle could also be transmitted as
10 bitmaps.

As generally represented in FIG. 2, a network 200, which may be wired, wireless or some combination thereof, includes at least two distinct users, each having an annotation device such as a tablet personal computer. In the example of FIG. 2,
15 two such devices 202₁ and 202₂ are shown, each connected to an annotations server 204, although a peer-to-peer model is one alternative. Note that the present invention has been designed to operate in a wireless (unreliable and limited bandwidth) network, as well as in a wired environment.

20 In general, as described below with reference to FIG. 3, the server 204 maintains associations between client publishers (authors) and client subscribers, and also collects the annotation data (e.g., in the form of strokes) from client authors, and transmits that data to client subscribers. Note

that although FIG. 2 shows two example devices 202₁ and 202₂, it should be understood that the present invention is intended to be able to work with many more than two such devices, essentially any practical number, e.g., limited only by the server capabilities and available bandwidth. For purposes of simplicity, in general one annotation device will be described herein, as the other annotation device includes the same components with respect to annotation capabilities (and thus each device's components differ only by the subscript, which ordinarily will be omitted herein).

Each annotation device 202 includes a client program 206, such as implemented in a custom browser component that uses the Internet Explorer control. The client program 206 launches local pages that are static, and pages that are dynamically obtained from a server, and hosts the shared inking functionality of the present invention, such as implemented in the form of dynamic link libraries (DLLs). The graphics and animations which present the user interface on the display (which in a tablet is also the input device) 208 may be implemented in C# using PNG (Portable Network Graphics) files with regions of low alpha blending used to increase the size of hit areas (tablet styli can be problematic for users to target accurately) or a custom application which manipulates layers. Note that one implementation of the

present invention has been designed to operate on various types of displays, including devices with lower graphic capabilities.

In accordance with an aspect of the present invention, a user is able to enter various modes of the annotations program 206, including a private mode in which users can see others' public annotations on the display 208 but no other users can see the private ink annotations data 210. Such a mode, which can be selected via buttons 402 and 410 (FIG. 4) or the like, is generally presented to the user in a manner that indicates that the user is not currently public, such as in the display 400 of FIG. 4. In this mode, a user can see the ink annotations of other users, (such as the writing "Hello Joe" in FIG. 4, but any private ink (such as the writing "Private ink") entered is not sent to the server 204. Note that although not readily apparent from FIG. 4, it should be understood that different color ink displays may be provided to differentiate a user's own ink (e.g., black) from that of others (e.g., blue), and/or differentiate the ink of others (e.g., blue for one other user's notes, green for another's notes, and so forth). A default color scheme may be provided, as described below, and the color scheme may also be user customizable.

As also represented in FIG. 4, other buttons 404 are available, including a "Subscribe to:" button for controlling to which others the client user is subscribing in order to see their ink annotations. Another of the buttons 404 is provided for obtaining (from the server) a list of any other users who have subscribed to see this user's public ink, and a third button in the button set 404 controls which annotations are presently to be displayed on the current user's display. Note that in one implementation, the ink annotations are on a per-page basis, and can accompany essentially any page of data, such as the graphic image 406 presented via a slideshow or the like. Note that the page of data may be a word processing document, a spreadsheet, and so on, (possibly allowing collaborative text editing), and that ink annotations may be added atop the publication page data as well as above, below and to the side of the page image. As shown in FIG. 2, the server 204 may provide this publication page data 212, e.g., as received from a publication device (e.g., running a presentation program) 214, in advance or dynamically, or the user may obtain it from another means, such as directly from the publication device 214 or possibly via a CD-ROM provided prior to a presentation. Note that the publication device 214 tells the server its current publication and page whenever changed, whereby the server 204 can notify the clients of a

change. Clients change publications and pages on demand, however it is feasible to have a model in which the server automatically controls the publication page that the clients see.

5 Controlling subscriptions to others' ink annotations is essentially performed by communicating with the server to obtain a list of other users (authors), as initiated via one of the buttons 404, e.g., the "Subscribe to:" button. One or more other authors may be selected from this list, such as via
10 a checkbox or the like, and the selection or selections returned to the server for associating any selected authors' identities with the subscribing user. A model in which a user can subscribe to his or her own published annotations is feasible, such as to see what those other subscribers are
15 seeing without private data also appearing. The list of available authors may be filtered in some way, such as to include only those users who have provided public data, however the list may include every participating user including those with no public annotations, thereby allowing a
20 user to subscribe in advance in the event that a user decides to later make some annotations public. Time-based filtering is also possible, e.g., show the private annotations along with the public annotations of user X and user Y as of a certain time.

A user can also use the buttons 404 (e.g., the "My Subscribers" button) to see which other users are subscribing thereto. Because the server maintains the associations on a per-user basis, this information is readily available from the server via a straightforward lookup. Note that although not necessary to the present invention, it is straightforward to provide a model in which a user can deny selected other users the ability to see his or her ink annotations, e.g., a semi-private mode of operation, either in advance or once the user realizes that another viewer is viewing his or her public annotations.

Another of the buttons 404 provides a "View Annotations" selection screen. This is essentially a filter mechanism, where the user can select "My Notes" and/or the notes of one or more other users for display. A checkbox on a window 408 may be used to provide an easy-to-use interface; (note that in the example of FIG. 4, the window 408 does not exactly correspond to the displayed ink, as there would be a selected checkbox for the source of the "Hello Joe" ink annotation) in order for that public ink of another to be displayed. In general, a user's own ink and the ink of others are maintained separately, and thus various layers of ink may be selectively overlaid into a resultant display.

To see the public ink annotations 220 (FIG. 2) of other users, in one implementation, a receive background thread 222 communicates with the server 204, identifying the user. In turn, the server returns the other users' published
5 annotations 220 to which that user has subscribed. This communication may be periodic, such as every few seconds, but may also be on demand to an extent and thus more immediate at certain times, such as whenever a user subscribes to a new author, or uses the view annotations interface checkbox to add
10 an existing author's annotations 220 to the current view. Buffering is performed as appropriate.

In accordance with another aspect of the present invention, another state is provided in the slide notations mode in which a user's notes are made public to other
15 subscribers. This mode may be entered via a public button 410. Although not separately shown, as described below with reference to FIG. 5, a number of indications are provided to a user whenever a user is in such a publishing state, to prevent a user from making a note that the user intends to be private
20 inadvertently become visible to others.

In general, entering this public state while in the slide notations mode makes the user's subsequent ink strokes (annotations) publicly available to subscribers. Once the user exits public mode, the user's published public

annotations persist on the other clients and the user's subsequent ink strokes (annotations) are once again private. As represented in FIG. 2, to this end, public ink annotations 230 are communicated to the server 204 via a background thread, referred to herein as a send thread 232 (although it is understood that the thread can perform other operations beyond data transmission, including collecting and buffering the public data to send). Note that such stroke data may also include data that erases previous stroke data. In general, the background thread gathers stroke data for public ink annotations and sends that ink to the server 304, such as by transmitting the data once every few seconds, and/or as public annotations data 230 becomes available. It should be understood that a separate background thread is only one way to provide the data, and that any other suitable mechanism is equivalent. In general, the data transmission typically identifies the user, identifies the user's current context (e.g., publication name and page), and provides the annotations in the form of stroke data to the server 204, for distribution to subscribers, as described below. As described below with reference to FIG. 3, the server tracks the public stroke data per user and context, such as in author and subscriber data structures 238.

In accordance with another aspect of the present invention, another useful annotations mode is a shared mode, also referred to as a shared canvas mode, as represented in FIG. 5 by the displayed page 500. In general, the shared
5 canvas mode allows multiple users (any user who wants) to see each others' notes that are written on a what appears to be a common display surface. In essence, every user who writes data to a shared canvas page 352_i-352_j (FIG. 3) is a published author, and every user is automatically a subscriber; the user
10 need only enter the shared canvas mode to participate. As with the public state when viewing a slide in the slide annotations mode, when in the shared canvas mode, a number of safeguards are provided so that the user is aware of the public state, including a visibly conspicuous (e.g., red and
15 white) graphic border 508, and a public state selector / indicator 510 being illuminated. Other indications such as flashing warning are equivalent. Note that in one implementation, the shared canvas mode is not strictly anonymous, since any user can choose to only show one layer of
20 ink at a time using the display filter, and thereby determine the source of a certain set of strokes. It is also feasible to have a bubble or the like pop up when hovering over ink to identify the source of that particular ink, and/or as ink is being added to the shared canvas.

Also, although not apparent in FIG. 4 or FIG. 5, a user's ink color may be made to automatically change when initially switching from public mode to private mode and vice-versa, although the user can manually change colors thereafter. For example, in one scheme, the user interface for the presentations portion uses the convention that the user's own ink is represented in black, though the user can change pen width or erase ink on a per stroke basis. When that user chooses to publish ink annotations by clicking on the public button, the ink becomes red as a reminder that what is being written is viewable by others. If subscribed to receive other users' ink, the ink may be represented in blue.

In general, the shared canvas mode is per-page corresponding to a page being presented, as controlled by the current presenter or a discussion moderator. As apparent from the example spontaneous, collaboration sketch / diagram shown in FIG. 5, such a shared canvas mode / page is highly useful to multiple user environments, such as an audience participating in a presentation or the like, as collaboration is automatic. A presenter can also find this highly useful, such as to conduct an informal survey, which is more anonymous than via a show of hands, and can provide more than simple yes no answers, e.g., users can mark an X where they believe they fit on a graph regarding some issue.

The presenter can clear the shared canvas by switching to another page, however that previous shared canvas may be persisted so that the shared notes for that page may reappear if the presenter goes back to the corresponding page. The
5 presenter can also erase the page as desired, again possibly after persisting one or more states of that shared canvas. Note that the server is associated with a non-volatile storage mechanism 240, and each annotation device has a non-volatile storage mechanism 242. Thus, annotations may be persisted and
10 recalled at the server 204 or a client 202 as needed, such as if the presenter goes back to a previous page. A user of an annotation device 202 can also operate offline and synchronize when online at a later time, and so on. To synchronize, a client author resends changed/new ink strokes and subscription
15 list changes made while offline to the server when back online, and if the server goes down, upon restart the server broadcasts its availability.

Yet another mode is a graffiti-type mode, represented via FIG. 6, in which users write to a canvas (graffiti page) 600
20 that is like the shared canvas, except that the graffiti page is not cleared, (at least not ordinarily), although a user can erase displayed ink. Thus, such a graffiti page is not per-page of a publication, but rather is a public data space (e.g., maintained on the server 350, FIG. 3) accessible to

every participant to add annotations and view the annotations of others, which may be per-publication, or per some other unit, such as the same page being available throughout an entire conference filled with multiple publications. Again, 5 some indication of the public state is provided, such as an illuminated indicator 610, a different background (bricks instead of lines) and so forth. The displayed wall may be scrolled, zoomed and so forth. Note that because ink in this mode is always public, a much greater variety of colors may be 10 provided since the colors scheme need not reserve colors for distinguishing between public versus private versus others' ink.

Turning to an explanation of the present invention with particular reference to FIG. 3, whenever a client subscribes 15 for receiving another users' public annotations, the server 204 associates that client subscriber with the annotation publisher (author) and his annotations for the specified document. For example, each publishing user may be listed in a table 310 (on a per-publication basis), with a pointer to a 20 per-author table 312 (or 313 and so on) of that author's subscribers. In this particular implementation, another table 314 (or 315 and so on) also having a pointer thereto from the table 310 contains the stroke data for that publication and author on a per-page basis.

Thus, whenever a background send thread provides published stroke data, the thread also provides the user identity and current context (publication and page), and the server adds the stroke data to the correct location in table 314. Note that shared canvas and graffiti data modes need not identify the user, but instead the server can simply merge the data into an appropriate set of shared stroke data 350 or 352 from multiple users, with shared canvas data being associated with a page (352_1-352_j), but not a user (in an anonymous model), and graffiti data simply added to a collection of strokes 350, which may be per publication, or multiple publications (e.g., for an entire conference having multiple presentations), or on any other suitable basis.

In one implementation, the present invention essentially operates as a push model, in which the server sends ink to subscribing clients' receive threads whenever changed ink is received from a publishing client. This is to avoid having the user wait for updates when switching to a page, and is particularly beneficial in a wireless environment, where the network may be disconnected or unreliable for any user.

For each author from which strokes are received, the author (e.g., Joe) can be found in the table 310, and the author's stroke data appropriately stored in the table 314,

and returned by the server to subscribers located via the table 312.

An alternative pull model is also possible, in which a user's receive thread requests published stroke data. In such a model, the receive thread identifies the requesting user, along with a publication, page and each author to which the user is currently subscribed (or at least seeking the public data thereof, e.g., to reduce network load, the user can only seek the data of those other users who are also selected for viewing rather than all authors to which that user has subscribed). The table 312 can be used to make sure that the requesting client has subscribed to the identified author. Users can provide a page and request shared canvas data, or can request graffiti data, and/or the server can update users as appropriate, e.g., by periodically sending deltas.

Note that in the alternative pull model, it is also straightforward to have the server track the authors for a subscriber, whereby the subscriber would not need to list the authors but instead would automatically receive the stroke data for each subscribed-to author. This would reduce the information that the receive thread had to provide with its request, however this would likely increase network load because stroke data would be returned for authors that the

client has subscribed to but may not be currently viewing the annotations thereof.

As described above (e.g., with reference to the "My Subscribers" button), one other set of data that a client may
5 receive in response to a request is a list of other user's that have subscribed to that client's public annotations. Thus, in the example of FIG. 3, the client author named "Joe" may request that the server return a list identifying his current subscribers, whereby the server 204 can access the
10 table 312 to find and return this information. In the example of FIG. 3, this would be Bill and Sam at the present moment. As can be readily appreciated, this information may be used in other ways as well, such as to allow any user to see another author's subscribers, to see which author was the most
15 subscribed-to author, and so on.

As can be seen from the foregoing detailed description, there is provided a method and system for sharing annotations that provides benefits like paper notes but further provides the advantages of computerized operations such as storage,
20 remote distribution and so forth. Annotations are selectively made publicly available as they are authored in a manner that handles one-to-many and many-to-many real time situations. The annotations are useable with any publication.

While the invention is susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions, and equivalents falling within the spirit and scope of the invention.